U.S. DEPARTMENT OF COMMERCE
National Technical Information Service

AD-A031 413

A Renewal Decision Problem

California Univ Berkeley Operations Research Center

Sep 76

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: ORC 76-28	2. GOVT ACCESSION NO.
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle): A RENEWAL DECISION PROBLEM	5. TYPE OF REPORT & PERIOD COVERED: Research Report
	6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s): C. Derman, G. J. Lieberman and S. M. Ross	8. CONTRACT OR GRANT NUMBER(s): DAAG29-76-G-0042
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Operations Research Center, University of California, Berkeley, California 94720	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: P-13179-M
11. CONTROLLING OFFICE NAME AND ADDRESS: U. S. Army Research Office, P.O. Box 12211, Research Triangle Park, North Carolina 27709	12. REPORT DATE: September 1976
	13. NUMBER OF PAGES: 21
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)	15. SECURITY CLASS. (of this report): Unclassified
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE: NA

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

NA

18. SUPPLEMENTARY NOTES

Also supported by the U. S. Army Research Office - Research Triangle Park under Grant DAHC04-75-G-0163.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Spares
Optimality Equation
Dynamic Programming
Random Termination
Fixed Termination
Discounted Cost
Replacements
Components
Component Category

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

(SEE ABSTRACT)

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

308133

ORC 76-28
SEPTEMBER 1976

A RENEWAL DECISION PROBLEM

by

C. DERMAN, G. J. LIEBERMAN

and

S. M. ROSS

OPERATIONS
RESEARCH
CENTER

UNIVERSITY OF CALIFORNIA · BERKELEY

21

A RENEWAL DECISION PROBLEM[†]

Operations Research Center Research Report No. 76-28

C. Derman, G. J. Lieberman and S. M. Ross

September 1976

U. S. Army Research Office - Research Triangle Park

DAAG29-76-G-0042

Department of Industrial Engineering
and Operations Research
University of California, Berkeley

APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION UNLIMITED.

[†]Partially supported by the U. S. Army Research Office - Research Triangle Park under Grant DAHC04-75-G-0163 with the University of California. Reproduction in whole or in part is permitted for any purpose of the United States Government.

THE FINDINGS IN THIS REPORT ARE NOT TO BE CONSTRUED AS AN OFFICIAL DEPARTMENT OF THE ARMY POSITION, UNLESS SO DESIGNATED BY OTHER AUTHORIZED DOCUMENTS.

ABSTRACT

A system must operate for t units of time. A certain component is essential for its operation and must be replaced, when it fails, with a new component. The class of spare components is grouped into n categories with components of the ith category costing a positive amount C_i and functioning for an exponential length of time with rate λ_i. The main problem of interest is, for a given t, to assign the initial component and subsequent replacements from among the n categories of spare components so as to minimize the expected cost of providing an operative component for t units of time.

In Section 1 we show that when there are an infinite number of spares of each category, the optimal policy has a simple structure. Namely, the time axis can be divided up into n intervals, some of which may be vacuous, such that when a replacement decision has to be made it is optimal to select a spare from the category having the ith largest value of λC whenever the remaining time falls into the ith closest interval to the origin. In Section 2 we consider the situation where $n = 2$ and there is only a single spare of one category and an infinite number of the other. In Section 3 we consider the case where there is only a finite number of spares for certain of the categories under the assumption that a rebate is allowed for the component in use at the end of the problem. In Section 4 we allude to a generalization of the model in Section 1 allowing for discounting or for the possibility that the system may randomly terminate before the t units of time expire. An optimal policy has the same simple structure as in Section 1.

A RENEWAL DECISION PROBLEM

by

C. Derman, G. J. Lieberman and S. Ross

0. Statement of Problem

A system must operate for t units of time. A certain component is essential for its operation and must be replaced, when it fails, with a new component. The class of spare components is grouped into n categories with components of the ith category costing a positive amount C_i and functioning for an exponential length of time with rate λ_i. The main problem of interest is, for a given t, to assign the initial component and subsequent replacements from among the n categories of spare components so as to minimize the expected cost of providing an operative component for t units of time.

In Section 1 we show that when there are an infinite number of spares of each category, the optimal policy has a simple structure. Namely, the time axis can be divided up into n intervals, some of which may be vacuous, such that when a replacement decision has to be made it is optimal to select a spare from the category having the ith largest value of λC whenever the remaining time falls into the ith closest interval to the origin. In Section 2 we consider the situation where $n = 2$ and there is only a single spare of one category and an infinite number of the other. In Section 3 we consider the case where there is only a finite number of spares for certain of the categories under the assumption that a rebate is allowed for the component in use at the end of the problem. In Section 4 we allude to a generalization of the model in Section 1 allowing for discounting or for the possibility that the system may randomly terminate before the t units of time expire. An optimal policy has the same simple structure as in Section 1.

1. Infinite Surplus in All Categories

In this section we suppose that our surplus of spare parts contains an infinite number of each category, and we number them so that $\lambda_1 C_1 > \lambda_2 C_2 > \ldots > \lambda_n C_n$. In addition, we suppose that there is no i and j such that $C_j \geq C_i$ and $\lambda_j \geq \lambda_i$; for if such is the case it can be shown (see Proposition 2) that category j need never be used.

Letting $V(t)$ denote the infimal expected additional cost incurred when there are t time units to go and a failure has just occurred, then $V(t)$ satisfies the optimality equation

$$V(t) = \min_{i=1,\ldots,n} \{C_i + \int_0^t V(t-x)\, \lambda_i e^{-\lambda_i x}\, dx\} , \quad t > 0 \qquad (1)$$

and $V(0) = 0$.

In addition, the policy which chooses, when t time units are remaining, a spare from a category whose number minimizes the right side of the optimality equation is an optimal policy. (This is a standard result in dynamic programming when all costs are assumed non-negative (see [3], [4]).

Proposition 1:

$V(t)$ is an increasing, continuous function of t for $t > 0$.

Proof: The increasing part follows from the definition of $V(t)$ since all costs are assumed non-negative. To prove continuity suppose that it is optimal to select a spare from category i whenever there are t units of time remaining. Then by selecting this same category at time $t + \varepsilon$ we see, from the lack of memory of the exponential, and the monotonicity of V that

$$V(t) \leq V(t + \varepsilon) \leq e^{-\lambda_i \varepsilon} V(t) + (1 - e^{-\lambda_i \varepsilon}) (C_i + V(t + \varepsilon)) .$$

Hence, the result is given. Q.E.D.

Theorem 1:

$V(t)$ is a piecewise linear concave function of t having at most n pieces.

Proof: Consider any value $t > 0$. Suppose the assignment of category i when t units of time remain is uniquely optimal. Then by the continuity of V and the optimality equation (1) there is an interval $(t, t + \varepsilon)$, $\varepsilon > 0$, such that i is uniquely optimal at every point in $(t, t + \varepsilon)$. Suppose several categories are optimal at t. Then the expressions within the brackets of (1) corresponding to each of the optimal categories are all equal to $V(t)$. If i is optimal the derivative of the expression with respect to t corresponding to category i is

$$\begin{aligned}
\frac{d}{dt}\{C_i + \int_0^t V(t - x)\, \lambda_i\, e^{-\lambda_i x}\, dx\} &= \frac{d}{dt}\{\int_0^t V(y)\, \lambda_i\, e^{-\lambda_i (t-y)}\, dy\} \\
&= \lambda_i\, V(t) - \lambda_i \int_0^t V(y)\, \lambda_i\, e^{-\lambda_i (t-y)}\, dy \\
&= \lambda_i\, V(t) - \lambda_i \int_0^t V(t - x)\, \lambda_i\, e^{-\lambda_i x}\, dx \qquad (2) \\
&= \lambda_i\, V(t) - \lambda_i (V(t) - C_i) \\
&= \lambda_i\, C_i\ ;
\end{aligned}$$

the derivative existing since $V(t)$ is continuous. It follows that among those categories that are optimal at t that category j with the smallest $\lambda_j C_j$ will be uniquely optimal over some interval $(t, t + \varepsilon')$, $\varepsilon' > 0$. Since at each change (as t increases) of optimal category a category with a smaller λC becomes optimal, there can be at most n values of t where a change in optimal category takes place. Since $\frac{dV(t)}{dt}$ is constant within the intervals where one category is optimal, $V(t)$ is linear within the interval; it is concave because the derivatives are non-increasing. Q.E.D.

Remark: Having established that an optimal policy employs the same category over intervals, the linearity and concavity of $V(t)$ can be deduced, as well, from the memoryless property of the exponential distribution and the linearity of the renewal function of the Poisson process.

It follows from the proof of Theorem 1 that the optimal policy uses category 1 spares when the time remaining is small, then switches to category 2 spares as the time increases, then category 3 spares as the time further increases, etc. where, of course, the interval of use for some categories may be empty. This suggests two possible algorithms for finding switching points.

Algorithm 1:

For this algorithm let V_i denote the minimal expected cost function and let π_i denote the optimal policy, when only categories $1, 2, \ldots, i$ are available. For instance

$$V_1(t) = C_1(1 + \lambda_1 t),\ 0 \leq t < \infty,$$

and π_1 is the policy which always replaces with a spare from category 1. From our previous structural results it follows that π_i will use category i whenever the time remaining is at least some finite critical value t_{i-1}. Now at t_{i-1}, it follows, by continuity, that it is optimal either to use category i and then proceed optimally, or to just use π_{i-1}. Hence, if $t_{i-1} > 0$, then

$$V_{i-1}(t_{i-1}) = C_i + \int_0^{t_{i-1}} V_{i-1}(t_{i-1} - x)\, \lambda_i e^{-\lambda_i x}\, dx\ . \quad (3)$$

Furthermore, since it follows from the optimality equation that for small values of t, π_i chooses the category with minimal value of C_k, we obtain

that $t_{i-1} = 0$ if and only if $C_i = \min_{1 \le k \le i} C_k$. Hence, unless this is the case, t_{i-1} can be taken to be the smallest positive solution of (3). And, in addition, we have

$$V_i(t) = \begin{cases} V_{i-1}(t) & t \le t_{i-1} \\ V_{i-1}(t_{i-1}) + \lambda_i C_i(t - t_{i-1}) & t \ge t_{i-1} \end{cases}$$

and π_i uses category i whenever $t \ge t_{i-1}$ and follows π_{i-1} when $t \le t_{i-1}$.

For example when $C_1 < C_2$, $\lambda_1 C_1 > \lambda_2 C_2$, this algorithm yields that t_1 is chosen so that

$$C_1(1 + \lambda_1 t_1) = C_2 + C_1 \int_0^{t_1} [1 + \lambda_1(t_1 - x)] \lambda_2 e^{-\lambda_2 x} dx .$$

Simplifying,

$$t_1 = \frac{1}{\lambda_2} \log \left[\frac{C_1\lambda_1 - C_1\lambda_2}{C_1\lambda_1 - C_2\lambda_2} \right] .$$

The expression for $V_2(t)$ can be written as

$$\begin{aligned} V_2(t) &= C_1(1 + \lambda_1 t) & t \le t_1 \\ &= C_1 + (\lambda_1 C_1 - \lambda_2 C_2) t_1 + \lambda_2 C_2 t & t \ge t_1 . \end{aligned}$$

<u>Algorithm 2</u>:

Let $C_{i_1} = \min\{C_1, \ldots, C_n\}$. For some value $t_1 > 0$, category i_1 is used whenever $0 < t \le t_1$. To find t_1, for every value of i, $i > i_1$, determine x_i, the smallest value of x, $x \ge 0$ satisfying

$$C_{i_1} + \int_0^x V(x - t)\, \lambda_{i_1} e^{-\lambda_{i_1} t}\, dt = C_i + \int_0^x V(x - t)\, \lambda_i e^{-\lambda_i t}\, dt$$

where

$$V(t) = C_{i_1}(\lambda_{i_1} t + 1) = u_1(t)\text{(say)} \quad \text{for} \quad t \geq 0 .$$

Let $t_1 = \min\limits_{i>i_1} \{x_i\} = x_{i_2}$. For some value t_2 category i_2 is used whenever $t_1 < t \leq t_2$. To find t_2, for every value of i, $i > i_2$, determine x_i, the smallest value of x, $x \geq t_1$ satisfying

$$C_{i_2} + \int_0^x V(x - t)\, \lambda_{i_2} e^{-\lambda_{i_2} t}\, dt = C_i + \int_0^x V(x - t)\, \lambda_i e^{-\lambda_i t}\, dt$$

where

$$V(t) = u_1(t) \quad \text{for} \quad 0 \leq t \leq t_1$$

$$= u_1(t_1) + \lambda_{i_2} C_{i_2}(t - t_1) = u_2(t) \text{ (say)} \quad \text{for} \quad t > t_1 .$$

Let $t_2 = \min\limits_{i>i_2} \{x_i\} = x_{i_3}$. For some value of t_3 category i_3 is used whenever $t_2 < t \leq t_3$. Recursively, category i_k is used whenever $t_{k-1} < t \leq t_k$ for some value t_k . To find t_k, for every value of i, $i > i_k$, determine x_i, the smallest value of x, $x \geq t_{k-1}$ satisfying

$$C_{i_k} + \int_0^x V(x - t)\, \lambda_{i_k} e^{-\lambda_{i_k} t}\, dt = C_i + \int_0^x V(x - t)\, \lambda_i e^{-\lambda_i t}\, dt$$

where

$$V(t) = u_j(t), \quad t_{j-1} < t \leq t_j, \quad j = 1,\ldots k-1$$

$$= u_{k-1}(t_{k-1}) + \lambda_{i_k} C_{i_k}(t - t_{k-1}) = u_k(t), \quad t > t_{k-1} .$$

Let $t_k = \min_{i>i_k} \{x_i\} = x_{i_{k+1}}$. This process stops when $t_k = \infty$. Of course if $i_k = n$, then $t_k = \infty$.

Both algorithms automatically exclude those categories which should never be used. However, it is possible to eliminate some in advance. This is indicated in the following.

<u>Proposition 2</u>.

If $C_j \geq C_i$, $\lambda_j > \lambda_i (C_j > C_i, \lambda_j \geq \lambda_i)$, then category j is never used in an optimal policy.

<u>Proof</u>: Let $t > 0$ be arbitrary. Let π_1 be the policy that uses category j at t and subsequently assigns categories optimally. Let π_2 be the policy that uses category i and subsequently assigns categories optimally. On comparing π_1 and π_2 we have

$$V_{\pi_1}(t) - V_{\pi_2}(t) = (C_j - C_i) + \left(\int_0^t V(t-x)\, \lambda_j\, e^{-\lambda_j x}\, dx - \int_0^t V(t-x)\, \lambda_i\, e^{-\lambda_i x}\, dx\right) .$$

The first expression on the right is non-negative (positive) by assumption. The second expression is positive (non-negative) since $\int_0^\infty f(x)\, \lambda\, e^{-\lambda x}\, dx$ is increasing in λ for every non-constant non-increasing function f ; $V(t - x)$ is such a function in x, as seen by letting $V(t - x) = 0$ for $x > t$. Thus $V_{\pi_1}(t) - V_{\pi_2}(t) > 0$. Since the use of category j can always be improved upon by using category i its use can never be optimal. Q.E.D.

<u>Remark</u>: It is also intuitive that if, for some i and j, $\lambda_i C_i < \lambda_j C_j$ and $C_i < C_j$, then category j will never be used. However, while this is evident from the formula for t_1 in the case of $n = 2$, and implies it is true for $n = 3$, we have not been able to prove it in general.

2. Finite Supply Model

In this section we suppose that $n = 2$ and that there is an infinite supply of spares of one category and only one of the other category.

Theorem 2:

If the set of spare components consists of one component of category 1 and an infinite number of category 2, where $C_2 > C_1$, then the optimal policy is to use a category 2 component when the time remaining t is greater than $\bar{t}$ and use the category 1 component when $t < \bar{t}$, where

$$\bar{t} = \frac{\log(C_2/C_1)}{\lambda_1-\lambda_2} \qquad \text{if} \qquad \lambda_1 > \lambda_2$$

$$= \infty \qquad \text{if} \qquad \lambda_1 \leq \lambda_2$$

Proof: Since once the decision to use the category 1 item is made there are no further decisions, it follows that one may regard this as a stopping rule problem where stopping means the use of the category 1 item. The one-stage look ahead stopping policy stops at t whenever stopping at t is better than continuing at t and then stopping at the next opportunity. Now letting X denote the lifetime of the category 1 component and V that of the first category 2 component used then W_1 the expected cost of using 1 at t, is

$$W_1 = C_1P\{X > t\} + (C_1 + C_2)\ P\{X < t,\ X + V > t\} + E[\text{cost}|X + V < t]\ \ P\{X + V < t\}\ ,$$

while W_2 the expected cost of first using a category 2 at t and then using the category 1 component, is given by

$$W_2 = C_2P\{V > t\} + (C_1 + C_2)\ P\{V < t,\ X + V > t\} + E[\text{cost}|X + V < t]\ P\{X + V < t\}\ .$$

Hence

$$W_1 - W_2 = C_1P\{X > t\} + (C_1 + C_2)\ P\{X < t,\ X + V > t\} - C_2P\{V > t\}$$

$$- (C_1 + C_2)\ P\{V < t,\ X + V > t\}$$

$$= C_1P\{V > t\} - C_2P\{X > t\}$$

$$= C_1\ e^{-\lambda_2 t} - C_2\ e^{-\lambda_1 t}\ .$$

Therefore, the one stage look ahead policy uses the category 1 component at t whenever

$$C_1\ e^{-\lambda_2 t} - C_2\ e^{-\lambda_1 t} \leq 0$$

or, equivalently, whenever

$$t \leq \frac{\log(C_2/C_1)}{\lambda_1 - \lambda_2} \qquad \text{if} \quad \lambda_1 > \lambda_2$$

or

$$t = \infty \qquad \text{if} \quad \lambda_1 \leq \lambda_2\ .$$

Since these sets of time points at which the one-stage look ahead policy stops can never be left once entered (without stopping), it follows (see [1] or [2]) that it is an optimal policy. Q.E.D.

<u>Remark</u>: The above proof does not need the assumption of exponential distributions for X and V . The same form of policy is optimal if there is a $\bar{t}$ such that

$$C_1\ P\{V > t\} - C_2\ P\{X > t\} \leq 0 \quad \text{for} \quad t \leq \bar{t}$$

$$(\geq 0) \qquad (t \geq \bar{t})$$

If the failure rate function of X is always greater than the failure rate function of V then such a $\bar{t}$ (possibly ∞) will exist since then $P(V > t)/P(X > t)$ is a non-decreasing function of t.

Theorem 3:

If the set of spare components is an infinite number of category 1 and one of category 2, where $C_2 > C_1$, $\lambda_2 < \lambda_1$, then the optimal policy is to use the category 2 component when the time remaining t is greater than x and use a category 1 component when $t \leq x$ where

$$x = \frac{1}{\lambda_2} \log \left[\frac{C_1\lambda_1 - C_1\lambda_2}{C_1\lambda_1 - C_2\lambda_2}\right]$$

Note: As we might expect x is the same as in the infinite supply model when $n = 2$.

Proof: Using the stated policy the expected cost function is

$$\begin{aligned} u(t) &= C_1(\lambda_1 t + 1) \quad \text{if} \quad t \leq x \\ &= C_2 + \int_0^t C_1[\lambda_1(t - y) + 1]\, \lambda_2 e^{-\lambda_2 y}\, dy \quad \text{if} \quad t > x . \end{aligned}$$

It is somewhat tedious but possible to verify that $u(t)$ satisfies

$$u(t) = \min\{C_1 + \int_0^t u(t - y)\, \lambda_1 e^{-\lambda_1 y}\, dy,\ C_2 + \int_0^t C_1[\lambda_1(t - y]\, \lambda_2 e^{-\lambda_2 y}\, dy\} .$$

That is, $u(t)$ satisfies the optimality equations for this problem; it thus follows from Proposition 1 of [3] that this policy is optimal. Q.E.D.

3. Finite Supply with Rebate Problem

Suppose in some of the n categories there are only a finite number of components. We assume that the components are numbered $j = 1, 2, \ldots$. However, in contrast to the previous problems, the cost of the last component used is returned. The problem, again, is to determine a policy for deciding which available component to use when a new component is required, so as to minimize the total expected cost.

If t units of time remain when a particular component from category i is put into use and L is the length of its life, then the expected cost associated with the use of this component is

$$
\begin{aligned}
E(\text{cost of component} \mid t) &= C_i P\{L < t\} \\
&= C_i \lambda_i \frac{1 - e^{-\lambda_i t}}{\lambda_i} \\
&= C_i \lambda_i \, E[\min(L,t)] \\
&= C_i \lambda_i \, E[\text{Length of time this component is used} \mid t] \; .
\end{aligned}
$$

Thus, letting $\delta(j)$ denote the category to which component j belongs it follows that the total expected cost under any policy π is

$$E_\pi[\text{total cost}] = \sum_j \lambda_{\delta(j)} C_{\delta(j)} E_\pi[\tau_j]$$

where τ_j is the amount of time component j is used, and where the sum is taken over all components.

Consider now a modified problem where the cost associated with component j is $\lambda_{\delta(j)} C_{\delta(j)} \tau_j$ where once again τ_j is defined as the amount of time that component j is in use, $j \geq 1$. The total expected cost with respect to the modified problem where policy π is used is precisely the same as the total expected cost with respect to the original problem. Thus, the policy that is optimal for the modified problem is optimal for the original problem. However, it is clear that the policy that used the category

associated with the minimal available $\lambda_i C_i$ is optimal for the modified problem. Thus we have proved the following:

<u>Theorem 4</u>:

The policy that minimizes the total expected cost when a rebate is given for the last component used is that policy that always selects among all the available categories that one having the smallest $\lambda_i C_i$.

<u>Remark 1</u>: The same problem can be thought of in a different context except that it leads to a maximization problem instead of a minimization one. There are n jobs to be performed sequentially within a fixed time t - the ith job takes an exponential amount of time with mean $1/\lambda_i$ and if completed within the time span of the problem earns the decision maker an amount C_i. Whenever a job is completed the decision maker must decide which job to attempt. He wishes to maximize the total expected earnings. The modified problem has the decision maker earn $\lambda_i C_i \tau_i$ if τ_i units of time are spent on the ith job whether or not it is completed within the time span of the problem. The decision maker wishes to maximize total expected earnings. As before, the optimal policy for the two problems is identical, namely, always choose the job with maximum $\lambda_i C_i$.

<u>Remark 2</u>: The conclusions for the problem in this section obviously hold if each category contains an infinite number of components. Since the optimal expected total cost must be less than the optimal expected cost for the problem discussed in Section 1 and since always using category n in that problem is not necessarily optimal we have the inequalities

$$\lambda_n C_n t \leq V(t) \leq \lambda_n C_n t + C_n .$$

Actually, better bounds can be obtained. If the rebate given for the last component used is $C_j - C_1$ when the last component used is from category j, it is still optimal to use category n for each replacement. One then arrives at

$$\lambda_n C_n t + C_1 \leq V(t).$$

<u>Remark 3</u>: Since the optimal policy is independent of the remaining time it follows that this policy is optimal when the time horizon is a random variable having any arbitrary distribution.

4. Discounted Case (A Random Termination Model)

Suppose in Section 1 discounting is appropriate, i.e., for some $\alpha > 0$

$$V(t) = \min_i \{C_i + \int_0^t e^{-\alpha x} V(t - x) \lambda_i e^{-\lambda_i x} dx\} .$$

This can be interpreted in the usual economic sense or as a random termination model. The latter interpretation arises when the system, in addition to being terminated definitely t units in the future, may also be terminated due to a randomly occurring accident; the time until such an accident occurs has an exponential distribution with mean $1/\alpha$.

The methodology of Section 1 applies in this case. Proposition 1 holds for $V(t)$. The derivative corresponding to (2) becomes $(\lambda_i + \alpha) C_i - \alpha V(t)$. With $\lambda_i C_i$ replaced by $(\lambda_i + \alpha) C_i$ the same type of statement concerning the structure of an optimal policy holds. The segments will no longer be piecewise linear; however, $V(t)$ will still be concave (the segments being the appropriate solution to the linear differential equation $V'(t) = (\lambda + \alpha) C - \alpha V(t))$.

REFERENCES

[1] CHOW, Y. S. and H. ROBBINS, "A Martingale Systems Theorem and Application:, Proc. Fourth Berkeley Symp. Math. Statistics Prob., Univ. of Calif. Press., 1961.

[2] DERMAN, C. and J. SACKS, "Replacement of Periodically Inspected Equipment (An optimal stopping rule), Naval Res. Logist. Quart., 1960, 597 - 607.

[3] ROSS, S. M., "Dynamic Programming and Gambling Models", Adv. in Appl. Prob., 1974, 593 - 606.

[4] STRAUCH, R., "Negative Dynamic Programming", Ann. Math. Statist., 1966, 171 - 189.